

REPORT DOCUMENTATION PAGE**Form Approved**
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 30-06-2009		2. REPORT TYPE Final		3. DATES COVERED (From - To) April 1, 2008 - March 31, 2009	
4. TITLE AND SUBTITLE Optimizing Machine Learning Algorithms for Hyperspectral Very Shallow Water (VSW) Products				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N00014-08-1-0622	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) W. Paul Bissett, Ph.D.				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Florida Environmental Research Institute PO Box 292397 Tampa, Florida 33687-2397				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research 875 North Randolph Street Arlington, VA 22203-1995				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release, Distribution is Unlimited					
13. SUPPLEMENTARY NOTES N/A					
14. ABSTRACT See Attached					
15. SUBJECT TERMS Inherent optical properties, IOP's, Hyperspectral Imagery (HSI), Message Passing Interface (MPI), Very Shallow Waters (VSW), bathymetry					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 2	19a. NAME OF RESPONSIBLE PERSON Beverly Walters
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code) 800-928-6402, ext 101

20090904429

This one-year effort focused on the transition of FERI's machine learning algorithms for HyperSpectral Imagery (HSI) in the Very Shallow Water (VSW) into a distributable code set. Our objective focused on two areas of application research and transitions. First, we transitioned our machine learning-based algorithms and computer code for the determination of bathymetry, bottom type, and water column Inherent Optical Properties from HyperSpectral Imagery (HSI) into a deliverable Message Passing Interface (MPI) code set that may be easily used by other research and military operators. Second, we moved beyond the use of single pixel HSI inversion to the use of spatial context-filtering to remove pixel-to-pixel noise inherent in the HSI data. In addition, the techniques and computer code used in this effort may be used with any set of spectral reflectance data, not just hyperspectral imagery. As such the deliverables from this effort will allow others to create maps of depths, bottom types, and water clarity from a variety of airborne and space-based spectral sensors planned for operational deployment.

Optimizing Machine Learning Algorithms For Hyperspectral Very Shallow Water (VSW) Products

W. Paul Bissett
Florida Environmental Research Institute
PO Box 292397
Tampa, FL 33687 USA

phone: (800) 928-6402 x102 fax: (800) 928-6402 email: pbissett [at] ferweb [dot] org

Award Number: N000140810622

<http://www.flenvironmental.org>

http://www.onr.navy.mil/sci_tech/32/322/ocean_optics_biology.asp

LONG-TERM GOALS

This one-year effort will focus on the transition of FERI's machine learning algorithms for HyperSpectral Imagery (HSI) in the VSW into a distributable code set. This will provide a stable code platform for the application and transition of machine learning-based hyperspectral classification techniques into 6.3/6.4 programs. (This work was funded mid-year 2008.)

OBJECTIVES

Our objective is to focus on three areas of application research and transitions. First, we will transition our machine learning-based algorithms and computer code for the determination of bathymetry, bottom type, and water column Inherent Optical Properties from HyperSpectral Imagery (HSI) into a deliverable Message Passing Interface (MPI) program that may be easily used by other research and military operators. Second, we will use this program to determine the impacts of the granularity of the classification database on the inversion bathymetry, bottom type, and IOPs. Third, we will move beyond the use of single pixel HSI inversion to the use of spatial context-filtering to remove pixel-to-pixel noise inherent in the HSI data.

APPROACH

Task 1

In previous works, a Look-Up Table (LUT) algorithm was used in accurately predicting bathymetry (Mobley et al. 2002, Bissett et al. 2004, Bissett et al. 2005, Mobley et al. 2005, Lesser and Mobley, 2008). The LUT approach is a subset of a larger body of artificial intelligence work concerned with algorithms and techniques that "teach" machines to learn from the examination of data and rules. This body of work is aptly called "machine learning" and some of its techniques include decision trees, genetic algorithms, and neural networks. The LUT approach is a subset of the k-Nearest Neighbor (kNN) algorithm, which is in the family of supervised learning algorithms.

Our use of the kNN algorithm maps a single HSI remote sensing reflectance vector, $Rrs(\lambda)$, onto a database of estimated $Rrs(\lambda)$. This database is created by providing the attributes of bathymetry, spectral bottom reflectance, and spectral IOPs to the radiative transfer routines of Ecolight (which is a high speed variant of Hydrolight, Mobley, 1994). We select the classification of the measured Rrs vector based on the best match of measured $Rrs(\lambda)$ to estimated $Rrs(\lambda)$. The LUT algorithm is based on a single best fit for our classification, i.e. $k = 1$. However, more recent work suggested that we could achieve a better classification by selecting a larger number for k , e.g. $k = 50$ (Bissett et al. 2006a). This larger number for k provides better accuracy and precision, as well as provides us with the ability to create confidence intervals for our classifications of bathymetry.

When classifying new spectra, the distance or angle between each measured spectrum and estimated spectrum in the database is calculated. The k nearest neighbors to that spectra (those having the smallest distances or angles), are considered sufficiently qualified to predict the corresponding attributes of bathymetry, bottom type, and IOP set. We have used the following metrics for the calculation of distance (Euclidean, Manhattan, Chebyshev, Canberra and Bray Curtis) and/or angle (Angular Separation and Correlation Coefficient). In general, our applications suggest that the Manhattan distance and the Correlation Coefficient angle metrics to be the best metrics to use for this algorithm. Once the set of nearest neighbors are determined, the attribute (e.g. bathymetry) of a pixel may be determined by a majority vote from the k nearest neighbor vectors. In the event of a tie, a prediction is made randomly from amongst the majority classes.

The computer code used in our creation of the estimated $Rrs(\lambda)$ database and the spectral matching of the measured versus estimated $Rrs(\lambda)$ is functional for scientific research; however it is not well developed for transition for use by others in testing and evaluation applications. **Our first task of this project will build upon our past research efforts to provide a Message Passing Interface (MPI) executable version of our kNN workbench for the inversion of hyperspectral imagery.** This code will be distributed to research and military partners for testing and evaluation purposes, as well as to complete Task 2 and 3.

Task 2

The spectrum for one particular depth, bottom type, and set of inherent optical properties may closely match a multitude of spectra with many different attributes (Figure 1). The selection of a single nearest neighbor may produce noisy predictions because of the noise in both the measured and estimated $Rrs(\lambda)$. The total prediction noise is a function of the noise associated with the measured $Rrs(\lambda)$, which contains components of sensor and environmental noise, and the noise associated with the estimation of $Rrs(\lambda)$ in the training database. This noise is evident in the "speckling" that may be associated with these inversion techniques (Figure 2). The use of kNN algorithms work to reduce noise of the prediction by increasing the probability that a spectrum presented for classification will come from the majority class of proximally-located spectral vectors, rather than a single "lucky" spectrum. In this case, rather than selecting the single database spectrum "O" that is closest to the measured spectrum (represented by the square in Figure 1), a majority vote of all of the nearest neighbors around the square is used to make the

prediction of the attribute (e.g. bathymetry) at that pixel location. Choosing the majority class creates a less variable space from which to make a decision, making it less likely to produce different classifications due to small amounts of noise in the spectra.

However, as the size of the training database increases (through the increase in number of bathymetry depths, bottom types, or IOP sets) the number of nearest neighbors also increases (Figure 3). This in turn causes a problem with “non-uniqueness” in the selection of the appropriate class, and its component attribute. This, in turn, causes increasing noise in the map of the estimated attribute (e.g. bathymetry), and therefore it becomes very important to have the appropriate “granularity”, or the proper step size in the discrete selection of attributes that are used in the creation of the training database. In this specific case, it means that we need to be selective in the selection of the number of depth levels, bottom types, and IOP sets that we use to create the estimated $Rrs(\lambda)$ database. **The second Task of this project will be to use the code from Task 1 to rapidly test the impacts of granularity of attribute selection on the accuracy and precision of bathymetry estimated from our kNN code and the HSI data from Horseshoe Reef and St. Joseph Bay, FL¹ (Bissett et al. 2006b).**

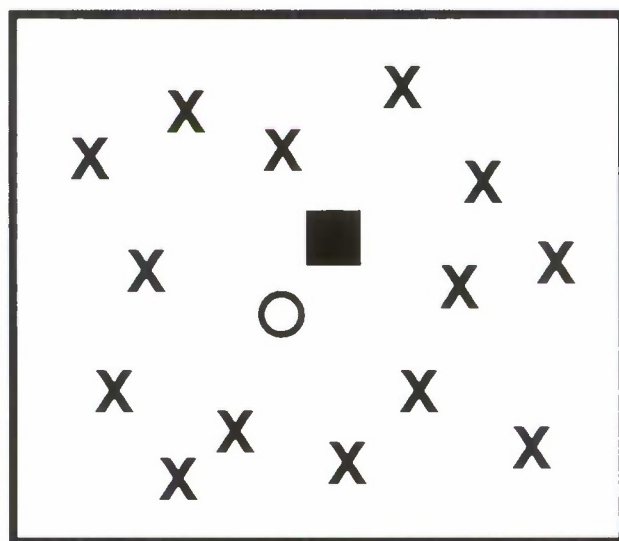


Figure 1. Xs and Os are the classes of examples belonging to the training database. The measured spectrum, \square , is closer to the O than any X. In kNN, multiple nearest neighbors are used to vote on the appropriate class. If $k = 1$, class O is chosen. If $k > 1$, a vote amongst all the classes X is chosen. The total number of Xs is dependent on the value of k , and which will include O in the retrieved set. The estimate of the attribute may then be calculated from any number of statistical calculations on the set of Xs, e.g. mean, majority vote, etc.

Task 3

The problem of sensor and environmental noise is a critical issue in the retrieval of accurate bathymetry from maps of HSI data. There are many sources of environmental noise in the

¹ The use of St. Joseph Bay, FL data will depend on acquiring accurate bathymetry from the State of Florida. If we do not receive bathymetry of sufficient quality, we will focus on the Horseshoe Reef imagery.

collection of sensor measured radiance, for example, surface waves that alter the reflection surface and path length to the bottom reflectance target. These surface noise effects are commingled with the atmospheric and illumination correction noise to produce spatially varying $R_{rs}(\lambda)$ over areas with identical bathymetry, bottom types, and IOPs (Figure 2). In order to reduce the impacts of this environmentally generated noise component, we should use the spatial context of the measured spectrum during the selection of the nearest neighbor classes, and subsequent estimate of the attribute of interest.

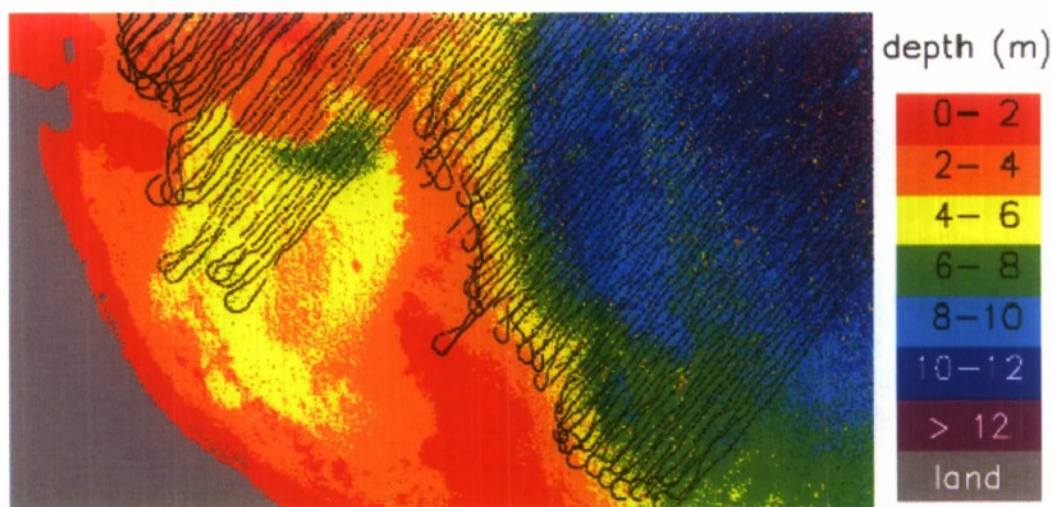


Figure 2. LUT bathymetry estimate for Horseshoe Reef, Bahamas. The black dots show the locations of the acoustic pings. The color-coded depths are for the unconstrained LUT retrieval ($k = 1$) applied to the entire image. The speckling in bathymetry is evident throughout the image.

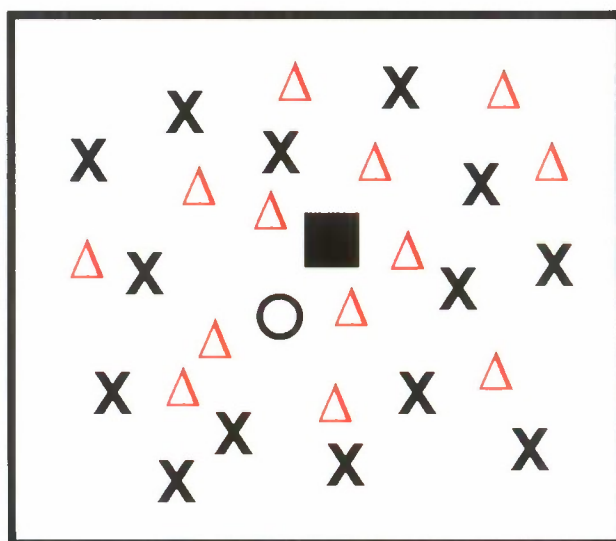


Figure 3. *Xs and Os are the classes of examples belonging to the training database and are the same as Figure 1. The Δ 's are additional classes resulting from increasing the depth resolution, as well as the number of bottom types and IOP sets. In this case discussed in the text, these Δ 's may contain attributes that are unrepresentative of the actual values and*

represents a non-unique solution to this inversion problem. The selection of the appropriate depth intervals or range of bottom types and IOPs sets is important to reducing this non-uniqueness. The term granularity is used to describe the separation between the discrete levels in the attributes.

Heretofore we have done point- or pixel-specific classification of HSI data. That is, each pixel is classified (for depth, bottom type, and water IOPs) independently of its neighbors, and only the spectral character of the pixel is used in its classification. Task 3 will be to evaluate spatial context-sensitive classification, which means that we will incorporate information about the spatial neighborhood (the spatial context) of a pixel to assist with its classification. Context-sensitive classification is often used in traditional terrestrial thematic mapping (e.g., Richards and Jia, 2006, §8.8) and some of those techniques may be beneficial for our oceanic problem.

This Task will evaluate two types of context-filtering – (1) pre-filtering of the $R_{rs}(\lambda)$ spectra before classification, and (2) context-filtering of the retrieved attributes after classification.

The first type of context-filtering seeks to reduce the noise in $R_{rs}(\lambda)$ spectra by replacing the spectrum value at each wavelength with the median value of the spectra in a spatial area surrounding the pixel of interest, say a 3×3 grid of pixels centered on the one of interest. This spatial filter is applied wavelength by wavelength. At wavelengths where $R_{rs}(\lambda)$ is mostly signal, the final spectrum will not change by much. At wavelengths where $R_{rs}(\lambda)$ is noisy, the noise in the surrounding pixels will tend to average out and the final spectrum values over the entire image area will be less noisy than the original.

The second type of context-filtering involves post-processing the retrievals themselves, rather than the original image spectra. In the case of real numbered attributes, such as bathymetry, we can apply a median filter to the retrieved depth. For bottom type and IOP set, the way forward is less clear. Each of these attributes is assigned a type with a specific vector (or set of vectors in the case of IOPs) of spectral values. How we filter “Dark Sediment” with “Sparse Vegetation” or “Highly absorbing and scattering waters #1” with “Case 1, chlorophyll $a = 0.5 \text{ mg m}^{-3}$ ” will be a challenge. It may require some iterative solution that context-filters bathymetry first, and solves the kNN again using a constrained bathymetry solution approach. It may also be highly dependent on the granularity study in Task 2. These are the issues that we will address in this Task.

WORK COMPLETED

Task (1) has been completed and the serial and MPI versions of our optimized machine learning code is available for v 0.1.0 release. The code will be distributed in a generic Red Hat Package Manager (RPM; http://en.wikipedia.org/wiki/RPM_Package_Manager) format for installation on Red Hat, Fedora, and CentOS version of Linux. This Task was expanded in anticipation of an ONR contract to transition this code set into an application appliance to be delivered to the Naval Oceanographic Office. This contract (N00014-09-C-0553) has been funded and the code set will be delivered September 2009.

The work for Task 2 and 3 starts with a baseline set of statistics with which to compare our spectral matching approaches to the “true” bathymetry measured with acoustical techniques. In

addition to previously used estimates (see below), we include a new estimation of “spikiness” in the retrieval of bathymetry from our spectrum matching techniques. Spikiness, S , is defined in the depth estimates as follows. For a given pixel (i,j) with retrieved depth $z(i,j)$, the average depth of the 4 neighboring pixels is

$$z_{avg4} = 0.25[z(i-1, j) + z(i+1, j) + z(i, j-1) + z(i, j+1)].$$

Spikiness, $S(i,j)$, of the retrieved depth at (i,j) as the absolute percent difference in depth $z(i,j)$ and z_{avg4} ,

$$S(i,j) = 100 \{ |z(i,j) - z_{avg4}| \} \text{ over } \{ z_{avg4} \}$$

For example at a $kNN=1$ (a single value LUT retrieval), if retrieval $z(i,j) = 5$ m or 15 m, and $z_{avg4} = 10$ m, then $S(i,j) = 50\%$. Note that a linearly sloping bottom is the same as a level bottom as regards the value of z_{avg4} . Thus a change in depth from one pixel to the next because of a sloping bottom is not recorded as spikiness. This metric is best suited for detecting a single spiky pixel. However, if a group of pixels is spiky, then some of the spiky pixels may be included in the z_{avg4} value, and the true spikiness may be underestimated for pixel (i,j) . Likewise, a sharp change in bottom depth, e.g., due to a coral head, may be recorded as a depth spike even though the LUT retrieval is correct.

Other statistical measures for “goodness of fit” from previous efforts include –

1. The average percent difference in LUT vs acoustic depths (a negative/positive value means that the LUT depths are on average shallower/deeper than the acoustic depths)
2. The average difference in meters in LUT vs acoustic depths (a negative/positive value means that the LUT depths are on average shallower/deeper than the acoustic depths)
3. The standard deviation in meters of the LUT vs acoustic depths
4. The correlation coefficient, r^2 , between the LUT and acoustic depths
5. The percent of pixels for which the LUT depth is within ± 1 m of the correct depth
6. The percent of pixels for which the LUT depth is within $\pm 25\%$ of the correct depth

The baseline for our comparison of various selections of spatial filtering parameters and kNN parameters is seen in Figures 4-6 and summarized in Figure 7. These figures show the bathymetry retrievals for unfiltered, $kNN = 1$ (LUT), parameters of our spectrum matching algorithms. In summary, we now have six quantitative measures of the overall accuracy of depth retrievals and two measures of the spikiness of depth retrievals. These metrics are used below to compare the effects of spatial smoothing of input Rrs spectra, of spatial smoothing of retrieved depths, and of the type of kNN analysis.

C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bl

no Rrs spatial smoothing
closest-match depth with no z spatial smoothing

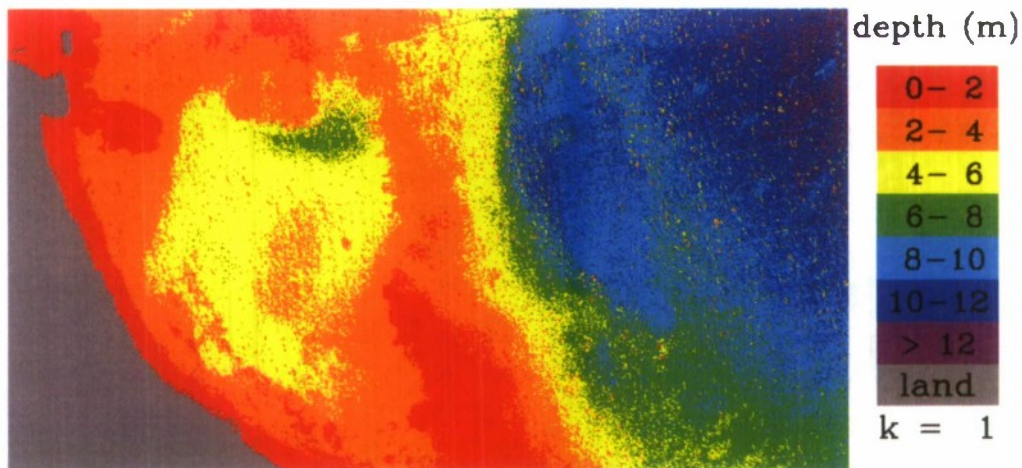


Figure 4. A 2D plot of retrieved depths, with the actual LUT-retrieved depths binned into 2-m bins and color-coded. Even with the binning, there is noticeable speckle in the deeper waters at the upper right.

C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bl

no Rrs spatial smoothing
closest-match depth with no z spatial smoothing

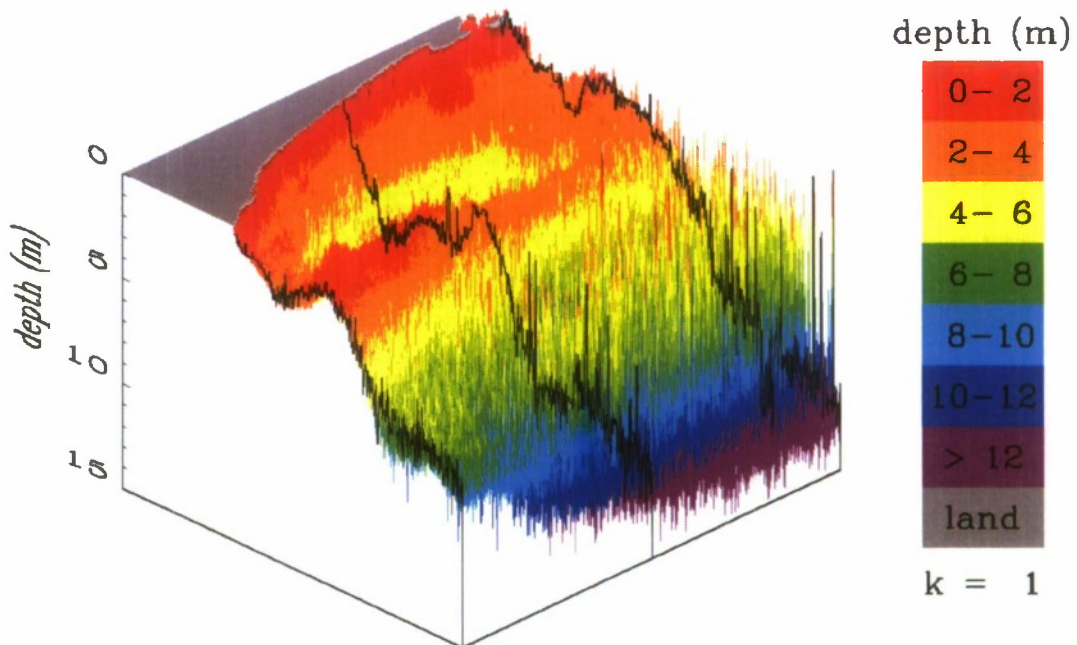


Figure 5. The LUT-retrieved depths plotted as a 3D surface and viewed in perspective (from the lower right direction of Fig. 4). The extreme spikiness of the depth retrievals is now quite apparent.

C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bl

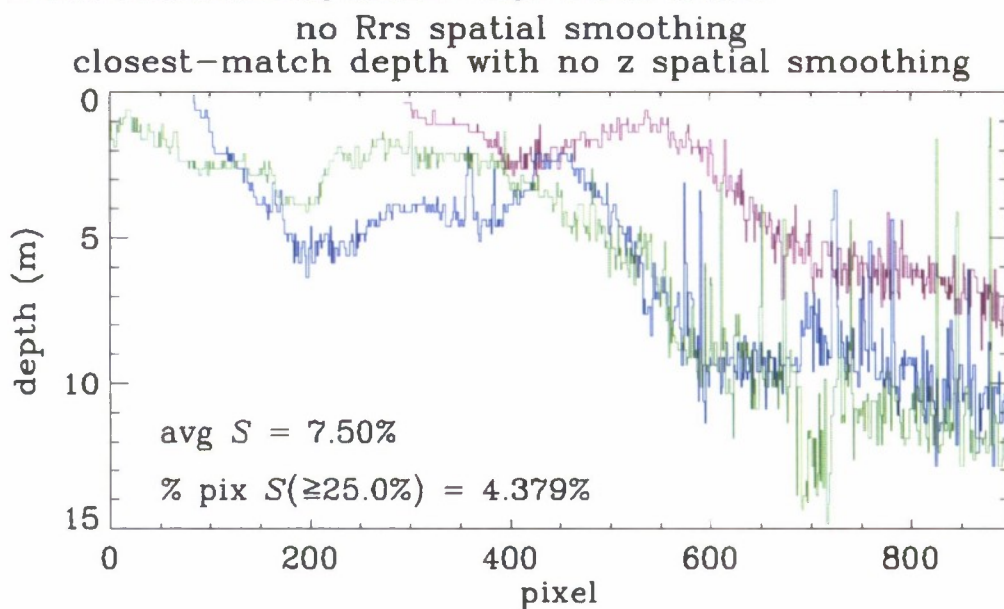
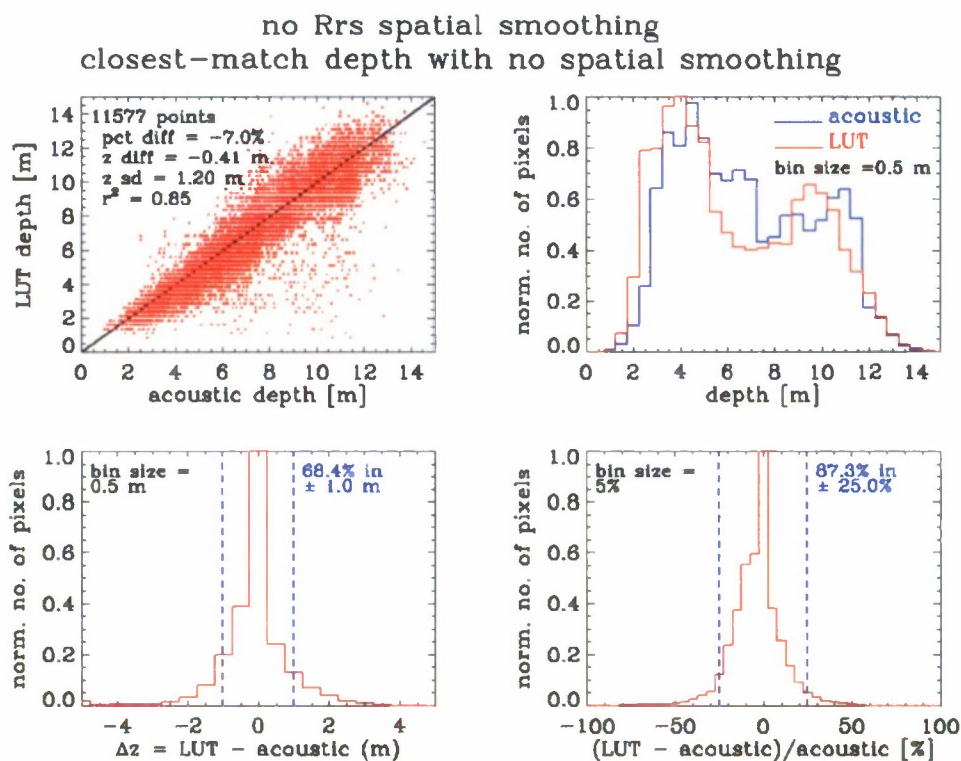


Figure 6. Depths along the 3 black transect lines seen in Fig. 5. The “bottom left” line in Fig. 5 is plotted in purple, the middle line in blue, and the “top right” line in green.



C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bl
c:\lut\phills\horseshoe\acoustic_bathymetry\comp_UTM_LL_HR2000_pix.txt

Figure 7. Goodness-of-fit results from LUT vs. acoustic depths for the baseline retrieval.

There are various ways of running a kNN algorithm to retrieve depth at each pixel that may impact the spikiness of the results, regardless of whether spatial smoothing of the Rrs or the retrieved attribute (e.g. depth) is performed. In performing the goodness-of-fit test, we wanted to consider the impact of differences in kNN selections in altering the results of the smoothing. The following are the basic criteria for kNN selections:

1. the closest match ($k = 1$, LUT)
2. the average of the $k = 30$ depths
3. the median of the $k = 30$ depths

It should be noted that kNN analysis does not reduce the retrieval error for pixels having whitecap or glitter contamination—if you start with a bad spectrum you get a bad result, no matter what the technique.

To spatially smooth an Rrs spectrum, we considered an $n \times n$ block of pixels centered on the pixel of interest, with $n = 1, 3$, and 5 ($n = 1$ corresponds to no spatial smoothing). Let $Rrs(i,j,\lambda)$ be the image spectrum at pixel (i,j) . To help eliminate anomalously large or small “bad” spectra, we discarded the highest and smallest values of the 9 spectra at each wavelength, and averaged the remaining 7 values. For $n = 5$, we discard the highest 2 and lowest 2 values, and averaged the remaining 21 values. If some of the pixels are flagged as land, clouds, or whitecaps, or if (i,j) is next to the image boundary, there are fewer than n^2 valid pixels, we discarded the highest and lowest values and average the remaining values. The original $Rrs(i,j,\lambda)$ is then replaced by the average spectrum computed from the $n \times n$ block of pixels. Note that this algorithm is applied independently at each wavelength. Thus, the particular spectra that are eliminated at one wavelength may or may not be the spectra that are eliminated at another wavelength.

To smooth the retrieved depths, we again consider $n \times n$ blocks of pixels. Now, however, we do not discard the high or low values of the retrieved depths before averaging. The reason is that when doing kNN matching, the kNN algorithm may have already omitted the high or low values, or done some other sort of filtering or averaging of the k retrieved depths at each pixel. We omit any pixels in the $n \times n$ block that are flagged as invalid (land, whitecap, image edge, etc), and then average the remaining (usually n^2) depths to obtain the spatially smoothed depth for the pixel at the center of the $n \times n$ block.

The matrix of combinations between kNN, Rrs, and depth averaging yield a $3 \times 3 \times 3$ solution matrix of 27 different combinations for analysis. That matrix and the results from the 8 tests are seen in Table 1. The following list provides a brief summary of the results.

R_{rs} smoothing	zb smoothing	kNN type	avg LUT vs % acous z % dif	avg LUT vs avg LUT vs acous z dif (m)	std dev of z dif (m)	% pix in +/- 1 m	% pix in +/- 25%	LUT vs acous r^2	avg spikiness S (%)	% pix with S > 25%	num pix with S > 25%
none	none	closest (k=1)	-7.0	-0.41	1.20	68.4	87.3	0.845	7.50	4.3800	16743
none	none	avg of k = 30	-1.8	-0.04	1.23	68.5	89.4	0.855	5.76	2.7900	10649
none	none	median of k = 30	-1.4	-0.01	1.26	67.2	88.7	0.851	6.23	3.4400	13143
none	3x3	closest (k=1)	-7.2	-0.43	0.88	75.3	91.9	0.909	1.49	0.0333	128
none	3x3	avg of k = 30	-1.9	-0.05	0.96	74.3	93.0	0.906	1.23	0.0279	107
none	3x3	median of k = 30	-1.5	-0.02	0.98	73.3	92.8	0.905	1.31	0.0404	155
none	5x5	closest (k=1)	-7.2	-0.43	0.80	77.1	93.4	0.924	0.69	0.0029	11
none	5x5	avg of k = 30	-2.0	-0.06	0.89	77.0	94.5	0.919	0.58	0.0016	6
none	5x5	median of k = 30	-1.6	-0.02	0.90	76.3	94.2	0.917	0.62	0.0052	20
3x3	none	closest (k=1)	-7.3	-0.43	0.93	73.4	90.7	0.901	4.39	1.3400	5123
3x3	none	avg of k = 30	-1.4	0.00	0.97	74.4	93.0	0.907	1.98	0.4254	1628
3x3	none	median of k = 30	-0.9	0.03	1.00	73.0	92.5	0.904	2.66	0.9085	3477
3x3	3x3	closest (k=1)	-7.4	-0.43	0.82	77.0	92.8	0.921	0.95	0.0193	74
3x3	3x3	avg of k = 30	-1.4	-0.01	0.92	75.9	93.9	0.917	0.66	0.0109	42
3x3	3x3	median of k = 30	-1.0	0.03	0.94	74.4	93.6	0.915	0.75	0.0182	70
3x3	5x5	closest (k=1)	-7.4	-0.44	0.78	78.3	93.6	0.929	0.45	0.0026	10
3x3	5x5	avg of k = 30	-1.5	-0.01	0.88	76.8	94.5	0.923	0.32	0.0005	2
3x3	5x5	median of k = 30	-1.0	0.02	0.90	75.8	94.3	0.921	0.36	0.0023	9
5x5	none	closest (k=1)	-7.3	-0.43	0.85	75.2	91.8	0.916	3.32	0.7966	3048
5x5	none	avg of k = 30	-1.3	0.01	0.90	76.5	94.1	0.920	1.25	0.1968	753
5x5	none	median of k = 30	-0.8	0.04	0.93	75.1	93.7	0.917	1.88	0.5567	2130
5x5	3x3	closest (k=1)	-7.4	-0.43	0.79	77.6	93.2	0.927	0.71	0.0096	37
5x5	3x3	avg of k = 30	-1.3	0.00	0.88	77.1	94.5	0.924	0.39	0.0029	11
5x5	3x3	median of k = 30	-0.8	0.04	0.91	75.7	94.3	0.921	0.49	0.0086	33
5x5	5x5	closest (k=1)	-7.4	-0.43	0.77	78.7	93.8	0.931	0.36	0.0013	5
5x5	5x5	avg of k = 30	-1.3	0.00	0.87	77.6	94.9	0.926	0.23	0.0010	4
5x5	5x5	median of k = 30	-0.8	0.04	0.89	76.2	94.6	0.923	0.27	0.0016	6

Table 1. Results of the 27 smoothing and kNN analyses.

1. kNN analysis does not help if the input Rrs spectrum is bad
2. Using the median of $k = 30$ depths gives slightly better signed depth errors than does the average of 30 depths
3. Using the average of $k = 30$ depths gives somewhat less spikiness (smaller average S values, and fewer pixels with $S > 25\%$) than does the median of $k = 30$ values
4. Other goodness-of-fit metrics are about the same for the average and median of $k = 30$ values
5. The average and median of $k = 30$ values give smaller signed depth errors (-0.8 to -2%) than does $k = 1$ (-7.0 to -7.4%), regardless of what smoothing is applied
6. The $k = 1$ depths give a smaller standard deviation of the LUT vs acoustic depth errors than does either the average or median of $k = 30$
7. Smoothing of the retrieved depths reduces spikiness much more than does a corresponding (having the same value of n) smoothing of the Rrs
8. The average of $k = 30$ values reduces both average and extreme spikiness more than does the median

These results are very encouraging when compared to our baseline retrievals (Figures 4 – 7). However, there is no single “best” methodology that gives superior values for all error metrics. Nevertheless, it appears that a reasonable recommendation (at least for the Horseshoe Reef image) is to:

1. use the median of $k = 30$ values to estimate the depth at each pixel (although using the average of $k = 30$ is about the same), which will give the most accurate average signed depth retrievals
2. definitely perform 3×3 or 5×5 spatial smoothing of the retrieved depths, which will greatly reduce the spikiness and thus further decrease the depth errors (Figure 8 - 11)
3. optionally also perform 3×3 or 5×5 spatial smoothing of the Rrs spectra before doing the LUT matching (Figure 12 - 15)

Figures 8 - 11 and 12 - 15 should be compared with Figures 4 - 7, which show the corresponding results for the baseline retrieval using no Rrs or z smoothing and $k = 1$ closest matching.

C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_15I-10P_Rb6-123_30NN.bil

no Rrs spatial smoothing

median of 30 depths with 3×3 z spatial smoothing

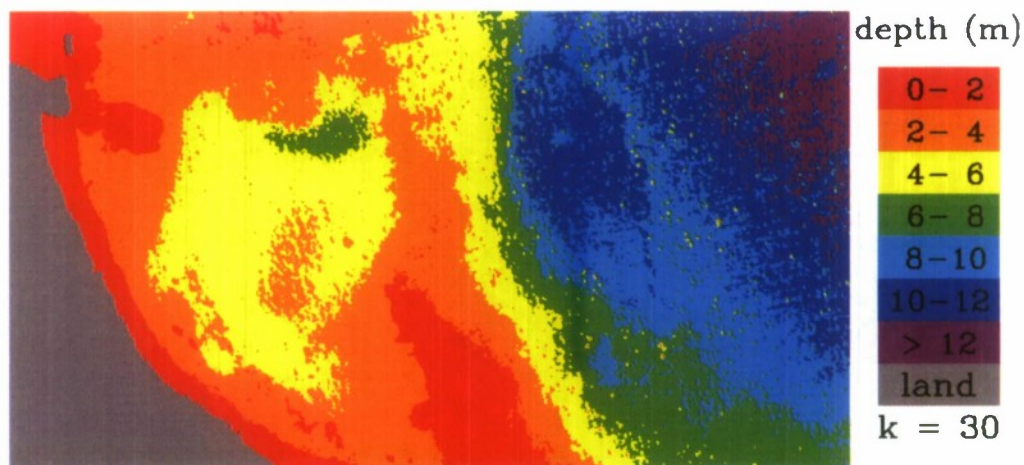


Figure 8. A 2D plot of retrieved depths, with the actual kNN-retrieved depths binned into 2-m bins and color-coded.

C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bil

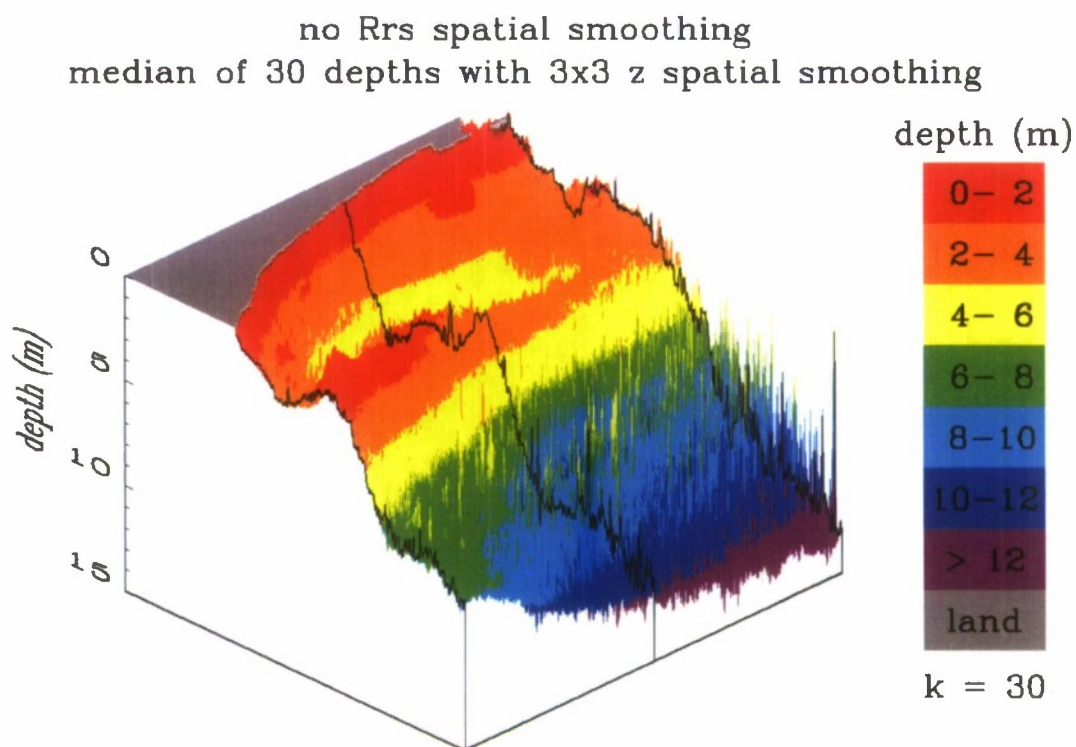


Figure 9. The LUT-retrieved depths plotted as a 3D surface and viewed in perspective (from the lower right direction of Fig. 8)

C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bil

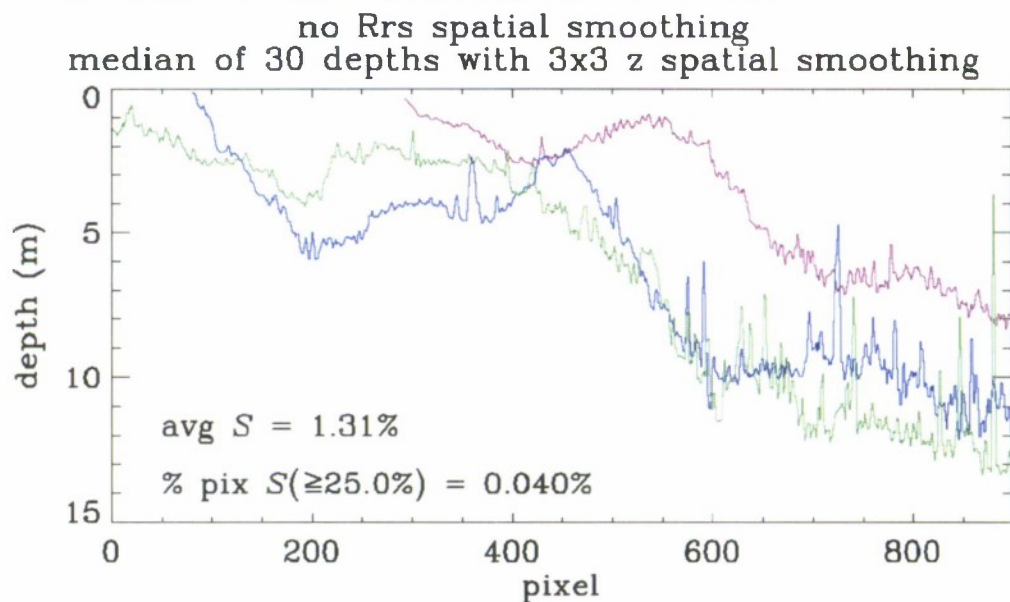
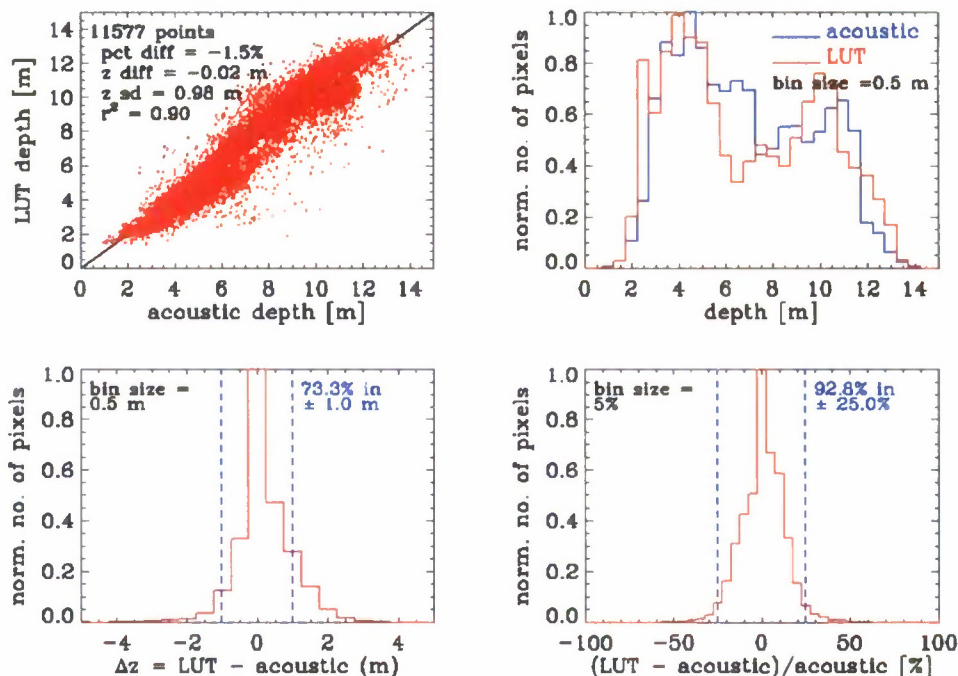


Figure 10. Depths along the 3 black transect lines seen in Fig. 9. The “bottom left” line in Fig. 9 is plotted in purple, the middle line in blue, and the “top right” line in green.

no Rrs spatial smoothing
median of 30 depths with 3x3 z spatial smoothing



C:\LUT\PHILLS\Horseshoe\HR2000_bathy_subsection_LUT_06Sep07_LSI-IOP_Rb6-123_30NN.bll
c:\lut\phills\horseshoe\acoustic_bathymetry\comp_UTM_LL_HR2000_pix.txt

Figure 11. Goodness-of-fit results from kNN vs. acoustic depths for the must-do retrieval.

C:\LUT\PHILLS\Horseshoe\HR2000_subsec_5x5midavg_LUT_03Apr08_LSI-IOP_Rb6-123_30NN.bll

5x5 Rrs spatial smoothing
median of 30 depths with 5x5 z spatial smoothing

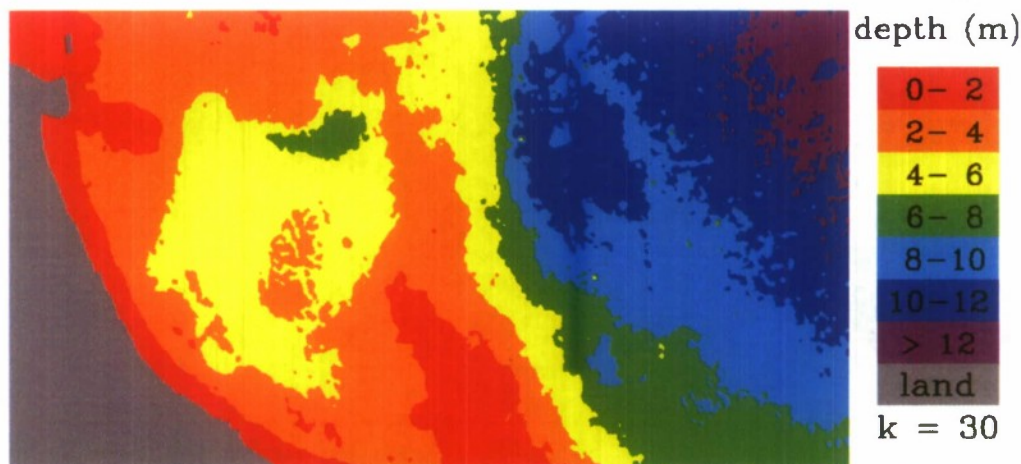


Figure 12. A 2D plot of retrieved depths, with the actual kNN-retrieved depths binned into 2-m bins and color-coded.

C:\LUT\PHILLS\Horseshoe\HR2000_subsec_5x5midavg_LUT_03Apr06_LSI-IOP_Rb6-123_30NN.bil

5x5 Rrs spatial smoothing
median of 30 depths with 5x5 z spatial smoothing

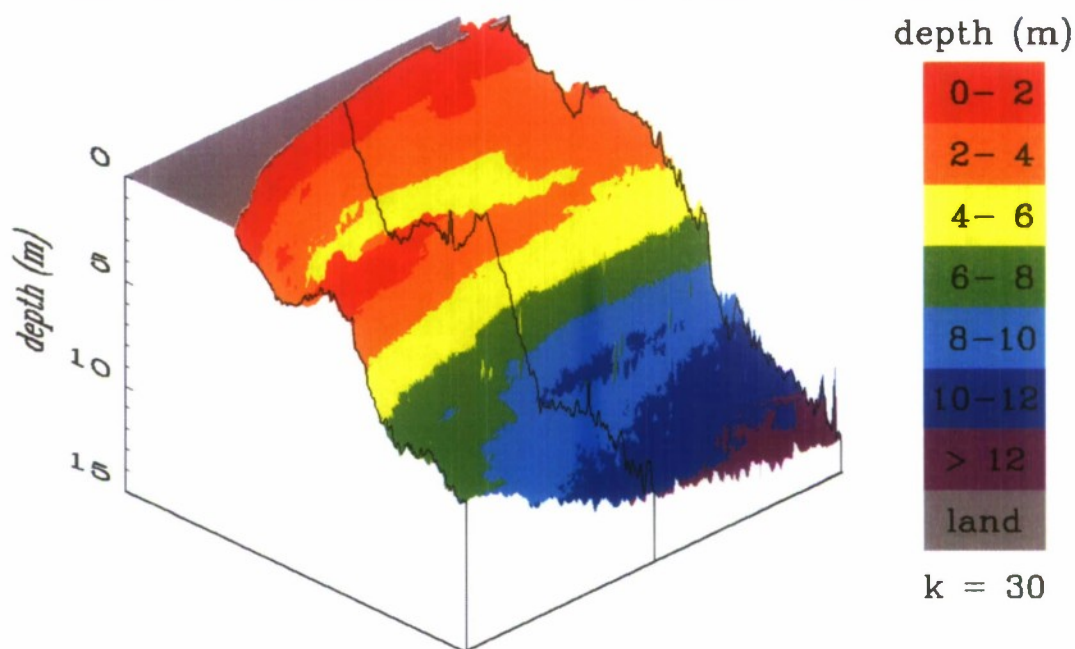


Figure 13. The LUT-retrieved depths plotted as a 3D surface and viewed in perspective (from the lower right direction of Fig. 12).

C:\LUT\PHILLS\Horseshoe\HR2000_subsec_5x5midavg_LUT_03Apr06_LSI-IOP_Rb6-123_30NN.bil

5x5 Rrs spatial smoothing
median of 30 depths with 5x5 z spatial smoothing

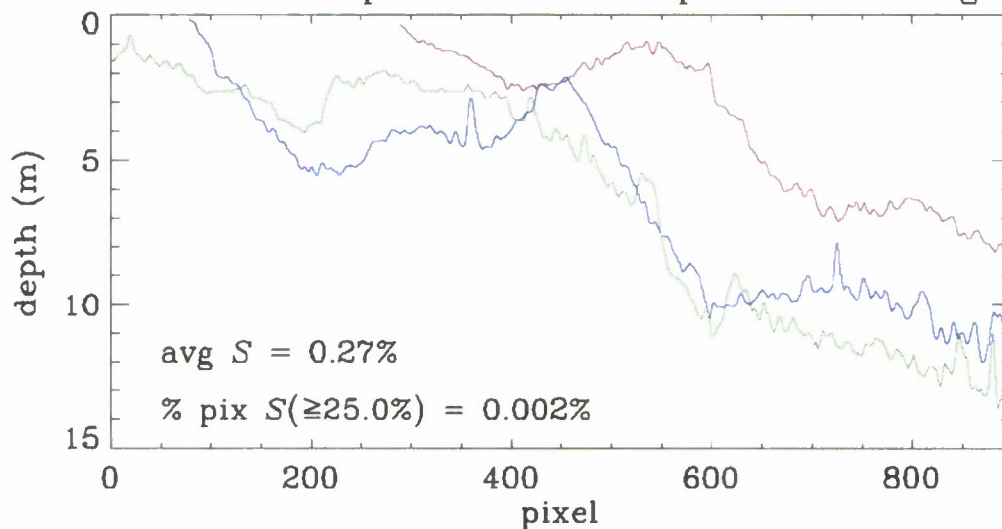
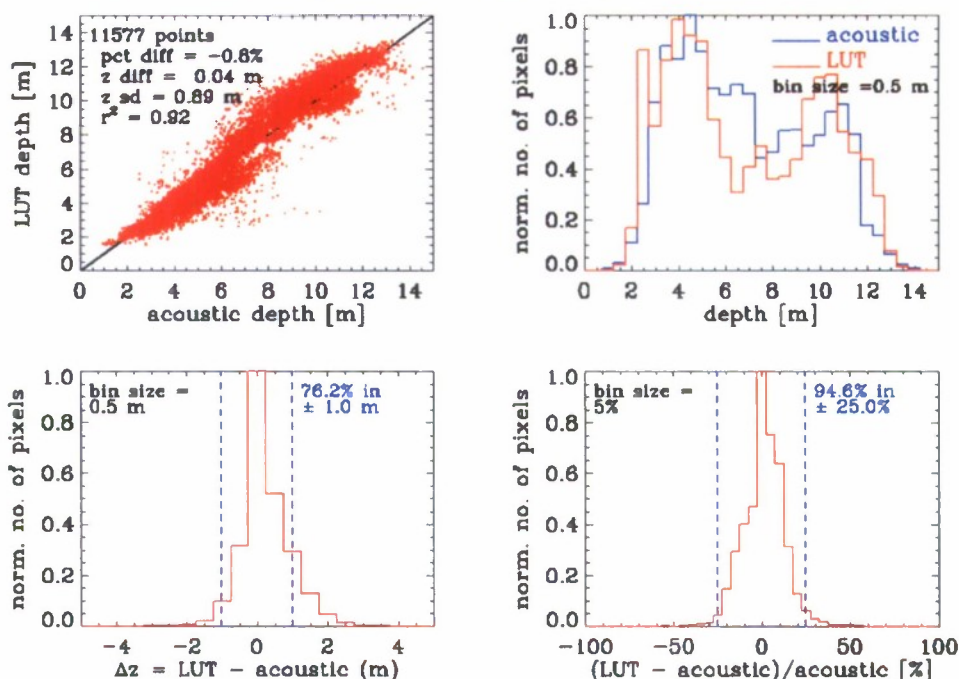


Figure 14. Depths along the 3 black transect lines seen in Fig. 13. The “bottom left” line in Fig. 13 is plotted in purple, the middle line in blue, and the “top right” line in green.

**5x5 Rrs spatial smoothing
median of 30 depths with 5x5 z spatial smoothing**



C:\LUT\PHILLS\Horseshoe\HR2000_subsec_5x5midavg_LUT_03Apr08_LSI-IOP_Rb6-123_30NN.bl
c:\lut\phills\horseshoe\acoustic_bathymetry\comp_UTM_LL_HR2000_pix.txt

Figure 15. Goodness-of-fit results from kNN vs. acoustic depths for optional retrieval.

IMPACT/APPLICATIONS

This effort will deliver an application for testing and evaluating of our machine learning approaches to bathymetry estimation in Very Shallow Waters (VSW). While it is being demonstrated on hyperspectral imagery, the techniques and computer code may be used with any set of spectral reflectance data. As such the deliverables from this effort will allow others to create maps of depths, bottom types, and water clarity from a variety of airborne and space-based spectral sensors planned for operational deployment.

RELATED PROJECTS

This work is being conducted in conjunction with Dr. Curtis D. Mobley at Sequoia Scientific, Inc., who is funded under this effort for the collaboration as well as under other collaborative spectrum matching funding. These techniques developed here are now being applied to imagery of Australian coastal waters in a comparison of several different hyperspectral remote sensing algorithms for a variety of environments. That comparison study is being led by A. Dekker of CSIRO. The kNN algorithms developed under this grant are being transitioned within an application appliance to be delivered to the Naval Oceanographic Office (N00014-09-C-0553) and is to be delivered September, 2009.

REFERENCES

Bissett, W.P., DeBra, S., Kadiwala, M., Kohler, D., Mobley, C., Steward, R., Weidemann, A., Davis, C.O., Lillycrop, J. and Pope, R., 2004. Development, validation, and fusion of high resolution active and passive optical imagery. Ocean Optics XVII, Fremantle, AU.

Bissett, W.P., DeBra, S., Kadiwala, M., Kohler, D.D.R., Mobley, C.D., Steward, R.G., Weidemann, A.D., Davis, C.O., Lillycrop, J. and Pope, R.L., 2005. Development, validation, and fusion of high-resolution active and passive optical imagery. In: I. Kadar (Editor), Signal Processing, Sensor Fusion, and Target Recognition XIV. Proceedings of SPIE Vol. 5809. SPIE, Bellingham, WA, pp. 341-349.

Bissett, W.P., Banfield, R., Kohler, D.D.R. and Mobley, C.D., 2006a. Ascribing Confidence Intervals to HyperSpectral Imaging (HSI) Bathymetry Maps, Ocean Optics XVIII, Montreal, Quebec, Canada.

Bissett, W.P. and Kohler, D.D.R., 2006. St. Joseph Bay Aquatic Preserve Hyperspectral Imaging - FINAL REPORT, Florida Environmental Research Institute, Tampa, FL. Florida Department of Environmental Protection, Contract Number RM055 (<http://www.feriweb.org/pubs/index.html#tech>)

Lesser, M. P. and C. D. Mobley. 2007. Bathymetry, optical properties, and benthic classification of coral reefs using hyperspectral remote sensing imagery. *Coral Reefs*, (26) 819-829.

Mobley, C.D., 1994. Light and Water. Academic Press, San Diego, CA, 592 pp.

Mobley, C.D., Sundman, L., Davis, C.O., Montes, M. and Bissett, W.P., 2002. A look-up-table approach to inverting remotely sensed ocean color data, Ocean Optics XVI. Office of Naval Research Ocean, Atmosphere, and Space S&T Department, Santa Fe, NM.

Mobley, C. D., L. K. Sundman, C. O. Davis, T. V. Downes, R. A. Leathers, M. J. Montes, J. H. Bowles, W. P. Bissett, D. D. R. Kohler, R. P. Reid, E. M. Louchard, and A. Gleason, 2005. Interpretation of hyperspectral remote-sensing imagery via spectrum matching and look-up tables, *Appl. Optics*, 44(17) 3576-3592.

Richards, J. A. and X. Jia, 2006. *Remote Sensing Digital Image Analysis, 4th Edition*. Springer, 439 pages.